

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Letters Patent of:
Thomas E. Kee et al.

Patent No.: 7,113,962

Issued: September 26, 2006

For: METHOD AND SYSTEM FOR
AUTOMATICALLY UPDATING CONTENT
STORED ON SERVERS CONNECTED BY A
NETWORK

**REQUEST FOR CERTIFICATE OF CORRECTION
PURSUANT TO 37 CFR 1.323**

Attention: Certificate of Correction Branch
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Sir:

Upon reviewing the above-identified patent, Patentee noted several patent office errors which should be corrected.

In the Specification:

Column 7, Lines 29-33 Delete "The Primary global. in greater detail." and insert the same below "Primary global server" as a new paragraph on Line 29.

Column 12, Line 64, Delete "requesters" and insert - - requestors - -.

Column 13, Line 59, Delete "requester" and insert - - requestor - -.

Enclosed please find marked up copies of pages 11, 20, 21 of the specification.

The following errors were found in the application as filed by applicant. The errors now sought to be corrected are inadvertent typographical errors. The correction of which does not involve new matter or require reexamination.

Sheet 5 of 14 (Box 166) (Fig. 3C), Line 1, Delete "REENABLE" and insert
- - RE-ENABLE - -.

Column 7, Line 28, After "server" insert -- . --.

Column 10, Line 50, Delete "<D1, 1.1.>" and insert - - <D1, 1.1> - -.

Column 19, Line 28, Delete "computer implemented" and insert
- - computer-implemented - -.

Transmitted herewith is a proposed Certificate of Correction effecting such amendment. Patentee respectfully solicits the granting of the requested Certificate of Correction.

The Commissioner is authorized to charge any deficiency of up to \$300.00 or credit any excess in this fee to Deposit Account No. 04-0100. Payment of \$100.00 is enclosed herewith.

Dated: December 6, 2006

Respectfully submitted,

By 

Flynn Barrison

Registration No.: 53,970

DARBY & DARBY P.C.

P.O. Box 5257

New York, New York 10150-5257

(212) 527-7700

(212) 527-7701 (Fax)

Attorneys/Agents For Applicant

**UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION**

Page 1 of 1

PATENT NO. : 7,113,962
APPLICATION NO. : 10/056,821
ISSUE DATE : September 26, 2006
INVENTOR(S) : Thomas E. Kee et al.

It is certified that an error appears or errors appear in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

In the Drawings:

Sheet 5 of 14 (Box 166) (Fig. 3C), Line 1, Delete "REENABLE" and insert
-- RE-ENABLE --.

In the Specification:

Column 7, Line 28, After "server" insert -- , --.

Column 7, Lines 29-33 Delete "The Primary global in greater detail." and insert the same below "Primary global server" as a new paragraph on Line 29.

Column 10, Line 50, Delete "<D1, 1.1,>" and insert -- <D1, 1.1> --.

Column 12, Line 64, Delete "requesters" and insert -- requestors --.

Column 13, Line 59, Delete "requester" and insert -- requestor --.

Column 19, Line 28, Delete "computer implemented" and insert
-- computer-implemented --.

MAILING ADDRESS OF SENDER:
Flynn Barrison
DARBY & DARBY P.C.
P.O. Box 5257
New York, New York 10150-5257

same name is not on the Primary global server, the source file on the source server is identified as the current version. Also, when a named source file only exists on a Primary global server, this source file is not identified as a member of the current version of the set of source files.

5 The process flows to a block 140 where the Primary global server gains file level access (FTP connection) to each source server that includes a source file that is identified as different than the current version of that particular file in the versioned file tree on the Primary global server. Each identified source file is copied to a new version in the versioned file tree repository on the Primary global server.

10 The Primary global server calls a library, e.g., the Revision Control Engine (RCE), to store file level differences between the current and previous versions of each source files. A discussion of FIGURES 6A and 6B below presents the functionality of the versioned file tree repository in greater detail.

15 In another embodiment, another file access protocol may be employed to transfer information, e.g., files, messages and data, between the Primary, Secondary, source and content servers. This other protocol could use a single port to enable all of the functions of the present invention, such as enabling the Primary global server to control the operation of the Secondary global server.

20 The process moves to a block 142 where the Primary global server generates version delivery lists and a list of Secondary global servers and their respective paths. Also, the Primary global server generates a version change container for each Secondary global server that may include a reference value associated with the current version of the set of source files.

25 Turning to FIGURE 3B from FIGURE 3A, the process advances to a block 146 where the Primary global server archives (compresses) each version change container. A third party facility may be used to implement a tape archive (TAR) command to compress each version change container. The process moves to a block 148 where a copy of the archived version change container is encrypted and transmitted to each Secondary global server. To reduce any adverse impact on the bandwidth
30 capacity of the network, each version change container may be broken down into

may have no other elements that need to be generated. This is what is generally referred to as a static Web page. Such pages may be created and then distributed to content servers. They may be distributed from source servers to content servers using mechanisms of primary global server 102 as described previously.

5 In contrast to origin servers, cache servers cache content and send it to requestors upon request. The content may be a copy of content on other servers such as an origin server, a content publishing system (such as primary global server 102), or another cache. By caching frequently requested content, a cache server may greatly decrease the time between a request and a response to the request (known as response
10 latency). A client, such as client 830, may be unaware that the content server responding is a cache server rather than an origin server.

A cache server, such as cache server 810, may be coupled more directly to primary global server 102, or it may be coupled to global server 102 over a network, such as Internet 101. Cache servers may be arranged in arrays, such as cache servers
15 815_{1-N} and 820_{1-N}. A device connected to Internet 101, such as a domain name system (DNS) server (not shown), may receive domain name requests from a client or a local domain name system (LDNS) server. The device may direct the client to a cache server or an origin server by returning an IP address associated with the cache or origin server. The device may direct requests depending on network traffic, network topology,
20 capacity of servers, content requested, and a host of other load balancing metrics. Cache servers 815_{1-N} may be located in one geographical location or may be spread to many different locations. Cache servers may be coupled to Internet 101 through firewalls or other packet-filtering systems for security and/or other reasons.

Server array controller 825 makes cache servers 820_{1-N} appear to
25 requestors as a single unified cache. It does this by redirecting a message directed to it to one of cache servers 820_{1-N}. It may redirect such messages based on many metrics including load balancing metrics. It may perform network address translation (NAT) and/or other translations, such as port address translation (PAT), to maintain the appearance that it is a single cache. Each cache server in FIGURE 8 could be replaced

with a server array controller controlling an array of cache servers without departing from the spirit or scope of the invention.

When distributing a version of content, primary global server 102 may interact differently with cache servers than with origin servers. For example, primary
5 global server 102 may cause some origin servers to go offline, i.e., to stop servicing requests for content, while allowing other origin servers to remain online to service requests occurring during updating the offline origin servers. For example, primary global server 102 may cause half the origin servers to go offline, update content on them, cause them to go online, and then cause the other half to go offline while updating
10 content on them. During the update process, requestors are able to get information from the online origin servers.

When primary global server 102 distributes a version of content, it may change operation of the cache servers. For example, it may cause the cache servers to go into a check-for-new-content mode in which each cache server sends a request for
15 content to an origin server to determine if new content exists and, if so, to cache it. In typical operation, when a cache server does not have requested content, it requests the content from an origin server. It then caches the content together with a time to live (TTL) field that indicates how long the cached content may be used before the cache server must re-request it from an origin server. The content together with the TTL field
20 and other information may be stored in a record called a cache entry (also called simply an entry). When a requestor requests content that the TTL field indicates has gone stale (passed its time to live), the cache server re-requests the content from an origin server.

When requesting content from an origin server, a cache server may request that the content be sent only if it has changed since the last request to the origin
25 server. For example, the cache server may send a time stamp or version of the content that indicates when the content was created. When an origin server receives such a request, it may compare the time stamp or version received from the cache server to the time stamp or version of the corresponding content the origin server has stored. If the origin server's content is more recent, it sends the content to the requestor which then
30 replaces its content and stores another time stamp or version with the content. If the